
Springer Tracts in Advanced Robotics

Volume 142

Series Editors

Bruno Siciliano, Dipartimento di Ingegneria Elettrica e Tecnologie dell'Informazione, Università degli Studi di Napoli Federico II, Napoli, Italy
Oussama Khatib, Artificial Intelligence Laboratory, Department of Computer Science, Stanford University, Stanford, CA, USA

Advisory Editors

Nancy Amato, Computer Science & Engineering, Texas A&M University, College Station, TX, USA
Oliver Brock, Fakultät IV, TU Berlin, Berlin, Germany
Herman Bruyninckx, KU Leuven, Heverlee, Belgium
Wolfram Burgard, Institute of Computer Science, University of Freiburg, Freiburg, Baden-Württemberg, Germany
Raja Chatila, ISIR, Paris cedex 05, France
Francois Chaumette, IRISA/INRIA, Rennes, Ardenes, France
Wan Kyun Chung, Robotics Laboratory, Mechanical Engineering, POSTECH, Pohang, Korea (Republic of)
Peter Corke, Queensland University of Technology, Brisbane, QLD, Australia
Paolo Dario, LEM, Scuola Superiore Sant'Anna, Pisa, Italy
Alessandro De Luca, DIAGAR, Sapienza Università di Roma, Roma, Italy
Rüdiger Dillmann, Humanoids and Intelligence Systems Lab, KIT - Karlsruher Institut für Technologie, Karlsruhe, Germany
Ken Goldberg, University of California, Berkeley, CA, USA
John Hollerbach, School of Computing, University of Utah, Salt Lake, UT, USA
Lydia E. Kavraki, Department of Computer Science, Rice University, Houston, TX, USA
Vijay Kumar, School of Engineering and Applied Mechanics, University of Pennsylvania, Philadelphia, PA, USA
Bradley J. Nelson, Institute of Robotics and Intelligent Systems, ETH Zurich, Zürich, Switzerland
Frank Chongwoo Park, Mechanical Engineering Department, Seoul National University, Seoul, Korea (Republic of)
S. E. Salcudean, The University of British Columbia, Vancouver, BC, Canada
Roland Siegwart, LEE J205, ETH Zürich, Institute of Robotics & Autonomous Systems Lab, Zürich, Switzerland
Gaurav S. Sukhatme, Department of Computer Science, University of Southern California, Los Angeles, CA, USA

The Springer Tracts in Advanced Robotics (STAR) publish new developments and advances in the fields of robotics research, rapidly and informally but with a high quality. The intent is to cover all the technical contents, applications, and multidisciplinary aspects of robotics, embedded in the fields of Mechanical Engineering, Computer Science, Electrical Engineering, Mechatronics, Control, and Life Sciences, as well as the methodologies behind them. Within the scope of the series are monographs, lecture notes, selected contributions from specialized conferences and workshops, as well as selected PhD theses.

Special offer: For all clients with a print standing order we offer free access to the electronic volumes of the Series published in the current year.

Indexed by SCOPUS, DBLP, EI Compendex, zbMATH, SCImago.

All books published in the series are submitted for consideration in Web of Science.

More information about this series at <http://www.springer.com/series/5208>

Peter Corke

Robotic Vision

Fundamental Algorithms in MATLAB®

Peter Corke
School of Electrical Engineering
and Robotics
Queensland University of Technology
Brisbane, QLD, Australia

ISSN 1610-7438 ISSN 1610-742X (electronic)
Springer Tracts in Advanced Robotics
ISBN 978-3-030-79174-2 ISBN 978-3-030-79175-9 (eBook)
<https://doi.org/10.1007/978-3-030-79175-9>

© Springer Nature Switzerland AG 2022

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

*To my family Phillipa, Lucy and Madeline for their indulgence and support;
my parents Margaret and David for kindling my curiosity;
and to Lou Paul who planted the seed that became this book.*

Foreword

At the dawn of the century's third decade, robotics is reaching an elevated level of maturity and continues to benefit from the advances and innovations in its enabling technologies. These all are contributing to an unprecedented effort to bringing robots to human environment in hospitals and homes, factories and schools; in the field for robots fighting fires, making goods and products, picking fruits and watering the farmland, saving time and lives. Robots today hold the promise for making a considerable impact in a wide range of real-world applications from industrial manufacturing to healthcare, transportation, and exploration of the deep space and sea. Tomorrow, robots will become pervasive and touch upon many aspects of modern life.

The *Springer Tracts in Advanced Robotics (STAR)* is devoted to bringing to the research community the latest advances in the robotics field on the basis of their significance and quality. Through a wide and timely dissemination of critical research developments in robotics, our objective with this series is to promote more exchanges and collaborations among the researchers in the community and contribute to further advancements in this rapidly growing field.

This is a refined remake of the volume of the second edition of *Robotics, Vision and Control – Fundamental Algorithms in MATLAB®* by Peter Corke in 2017. The work now comes in two split volumes: one devoted to *Robotics and Control*, and the other to *Robotic Vision*. The first volume contains material from the first nine chapters of the previous single volume, covering: foundations on pose, time, and motion; mobile robots with navigation and localization; kinematics, dynamics, and control of robot manipulators. On the other hand, the second volume contains material from the first two chapters and the tenth to fourteenth chapters of the previous single volume, covering: foundations on pose, computer vision, image processing and feature extraction; image formation and multiple images for the geometry of vision.

The outcome is a two-volume handy set which is confirmed to be shining in our STAR series!

Naples, Italy and Stanford, USA
November 2020

Bruno Siciliano and Oussama Khatib
STAR Editors

Foreword to the Second Edition

Once upon a time, a very thick document of a dissertation from a faraway land came to me for evaluation. *Visual robot control* was the thesis theme and *Peter Corke* was its author. Here, I am reminded of an excerpt of my comments, which reads, *this is a masterful document, a quality of thesis one would like all of one's students to strive for, knowing very few could attain – very well considered and executed.*

The connection between robotics and vision has been, for over two decades, the central thread of Peter Corke's productive investigations and successful developments and implementations. This rare experience is bearing fruit in this second edition of his book on *Robotics, Vision, and Control*. In its melding of theory and application, this second edition has considerably benefited from the author's unique mix of academic and real-world application influences through his many years of work in robotic mining, flying, underwater, and field robotics.

There have been numerous textbooks in robotics and vision, but few have reached the level of integration, analysis, dissection, and practical illustrations evidenced in this book. The discussion is thorough, the narrative is remarkably informative and accessible, and the overall impression is of a significant contribution for researchers and future investigators in our field. Most every element that could be considered as relevant to the task seems to have been analyzed and incorporated, and the effective use of Toolbox software echoes this thoroughness.

The reader is taken on a realistic walkthrough the fundamentals of mobile robots, navigation, localization, manipulator-arm kinematics, dynamics, and joint-level control, as well as camera modeling, image processing, feature extraction, and multi-view geometry. These areas are finally brought together through extensive discussion of visual servo system. In the process, the author provides insights into how complex problems can be decomposed and solved using powerful numerical tools and effective software.

The *Springer Tracts in Advanced Robotics (STAR)* is devoted to bringing to the research community the latest advances in the robotics field on the basis of their significance and quality. Through a wide and timely dissemination of critical research developments in robotics, our objective with this series is to promote more exchanges and collaborations among the researchers in the community and contribute to further advancements in this rapidly growing field.

Peter Corke brings a great addition to our STAR series with an authoritative book, reaching across fields, thoughtfully conceived and brilliantly accomplished.

Oussama Khatib
Stanford, California
October 2016

Preface

*Tell me and I will forget.
Show me and I will remember.
Involve me and I will understand.*
Chinese proverb

*Simple things should be simple,
complex things should be possible.*
Alan Kay

Digital images are everywhere – most of us have collections of thousands of digital images captured on a variety of devices. Each image is rich in information about the scene it recorded – implicit in every pixel is information about the 3D structure of the scene, the materials which comprise it and the lighting conditions at the time. In our everyday life we use our sense of vision to help us perform a huge variety of tasks: recognizing people, reading text, manipulating objects and navigating. Some simple visual skills can now be performed by machines, for example, cameras can tell if we're smiling, social media platforms and customs barriers can recognize individual faces; and we can search our online photo albums using text keywords. Robots need mastery of visual skills if they are to meet their potential as helpful and ubiquitous machines in our everyday lives.

Robots are data-driven machines – they acquire data, process it and take action based on it. The data comes from many sensors including cameras. An image comprises millions of pixels and the amount of data that needs to be processed, in real-time, is massive.

Progress robotic vision has been, and continues to be, driven by more effective ways to process that data. This has been achieved through new and more efficient algorithms, as well as the dramatic increase in computational power that follows Moore's law.◀ When I started in robotics and vision in the mid 1980s, see Fig. 0.1, the IBM PC had been recently released – it had a 4.77 MHz 16-bit microprocessor and 16 kbytes (expandable to 256 k) of memory. Over the intervening 30+ years computing power has perhaps doubled 20 times which is an increase by a factor of over one million. In the last 5 years, graphical processing units, or GPUs, have come to the fore and powered a new wave of data-driven algorithms most notably deep learning◀.

Over the fairly recent history of robotic vision a very large body of algorithms has been developed to efficiently solve large-scale problems in scene understanding◀, motion estimation, and 3D reconstruction – a significant, tangible, and collective achievement of the research community. However its sheer size and complexity presents a very real barrier to somebody new entering the field. Given so many algorithms from which to choose, a real and important question is:

What is the right algorithm for my particular problem?

One strategy would be to try a few different algorithms and see which works best for the problem at hand, but this is not trivial and leads to the next question:

How can I evaluate algorithm X on my own data without spending days coding and debugging it from the original research papers?

Two developments come to our aid. The first is the availability of general purpose mathematical software which makes it easy to prototype algorithms. There are commercial packages such as MATLAB®, Mathematica®, Maple® and MathCad®◀, as well as open source projects such as SciLab, Octave, and Matplotlib. All these tools deal naturally and effortlessly with vectors and matrices, can create complex and beautiful

"Computers in the future may weigh no more than 1.5 tons." Popular Mechanics, forecasting the relentless march of science, 1949

The topic of deep learning is outside the scope of this book.

"Understanding" is a word with strong connotations and it is a philosophical question as to whether a machine can ever "understand". In this context, we use a weak interpretation of "understanding" which encompasses recognizing objects and perhaps activities.

Respectively the trademarks of The MathWorks Inc., Wolfram Research, MapleSoft and PTC.

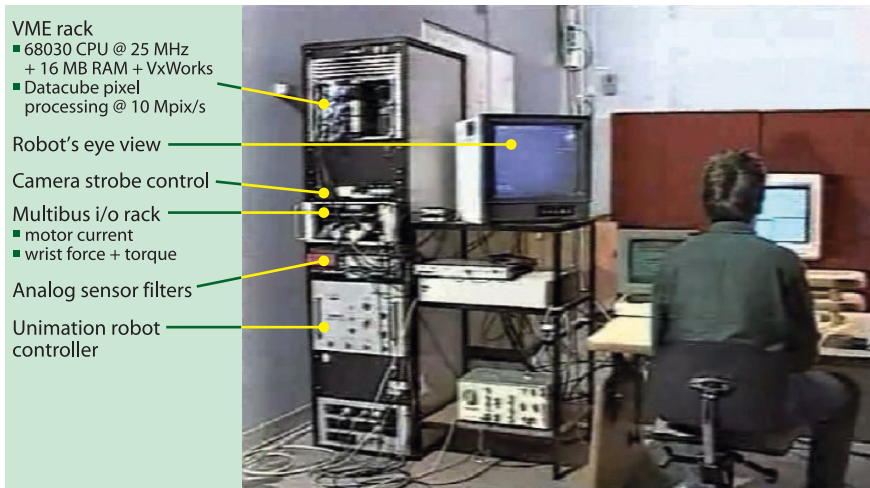


Fig. 0.1.

Once upon a time a lot of equipment was needed to do vision-based robot control. The author with a large rack full of real-time image processing and robot control equipment (1992)

graphics, and can be used interactively or as a programming environment. The second is the open-source movement. Many algorithms developed by researchers are available in open-source form. They might be coded in one of the general purpose mathematical languages just mentioned, or written in a mainstream language like C, C++, Java or Python.

For more than twenty years I have been part of the open-source community and maintained two open-source MATLAB Toolboxes: one for machine vision and one for robotics. They date back to my own Ph.D. work and have evolved since then, growing features and tracking changes to the MATLAB language. They form the basis of this book.

These Toolboxes have some important virtues. Firstly, they have been around for a long time and have been used by many people for many different problems so the code can be accorded some level of trust. New algorithms, or even the same algorithms coded in new languages or executing in new environments, can be compared against implementations in the Toolbox.

» allow the user to work with real problems,
not just trivial examples

Secondly, they allow the user to work with real problems, not just trivial examples. For real images with millions of pixels the computation required is beyond unaided human ability. Thirdly, they allow us to gain insight which can otherwise get lost in the complexity. We can rapidly and easily experiment, play *what if* games, and depict the results graphically using the powerful 2D and 3D graphical display tools of MATLAB. Fourthly, the Toolbox code makes many common algorithms tangible and accessible. You can read the code, you can apply it to your own problems, and you can extend it or rewrite it. It gives you a “head start” as you begin your journey into robotic vision.

The book takes a conversational approach, weaving text, mathematics and code examples into a narrative. I wanted to show how complex problems can be decomposed and solved using just a few simple lines of code. More formally this is an inductive learning approach, going from specific and concrete examples to the more general.

» show how complex problems can be decomposed
and solved

The topics covered in this book are based on my own interests but also guided by real problems that I observed over many years as a practitioner of both robotics and computer vision. I want to give you a flavor of what robotic vision is about and what

it can do – consider it a grand tasting menu. I hope that by the end of this book you will share my enthusiasm for these topics.

» consider it a grand tasting menu

I was particularly motivated to present a solid introduction to computer vision for roboticists. The treatment of vision in robotics textbooks tends to concentrate on simple binary vision techniques. In this book we will cover a broad range of topics including color vision, advanced segmentation techniques, image warping, stereo vision, motion estimation, bundle adjustment, visual odometry and image retrieval. We also cover nonperspective imaging using fisheye lenses, catadioptric optics and the emerging area of light-field cameras. These topics are growing in importance for robotics but are not commonly covered. Vision is a powerful sensor, and roboticists should have a solid grounding in modern fundamentals.

» software is a first-class citizen in this book

This book is unlike other text books, and deliberately so. Firstly, software is a first-class citizen in this book. Software is a tangible instantiation of the algorithms described – it can be read and it can be pulled apart, modified and put back together again. There are a number of classic books that use software in an illustrative fashion and have influenced my approach, for example *LaTeX: A document preparation system* (Lamport 1994), *Numerical Recipes in C* (Press et al. 2007), *The Little Lisper* (Friedman et al. 1987) and *Structure and Interpretation of Classical Mechanics* (Sussman et al. 2001). Over 700 examples in this book illustrate how the Toolbox software can be used and generally provide *instant gratification* in just a couple of lines of MATLAB code.

» instant gratification in just a couple of lines of MATLAB code

Secondly, building the book around MATLAB and the Toolbox means that we are able to tackle more realistic and more complex problems than other books.

» this book provides a complementary approach

The emphasis on software and examples does not mean that rigor and theory are unimportant – they are very important, but this book provides a complementary approach. It is best read in conjunction with other texts which do offer rigor and theoretical nourishment. The end of each chapter has a section on further reading and provides pointers to relevant textbooks, key papers and online resources. I try hard to use the least amount of mathematical notation required, but if you seek deep mathematical rigor this may not be the book for you.

The Toolbox also includes some great open-source software and I am grateful to the following for code that has been incorporated into the Machine Vision Toolbox: RANSAC by Peter Kovesi; pose estimation by Francesco Moreno-Noguer, Vincent Lepetit, and Pascal Fua; color space conversions by Pascal Getreuer; numerical routines for geometric vision by various members of the Visual Geometry Group at Oxford (from the web site of the book by Hartley and Zisserman 2003); *k*-means, SIFT and MSER algorithms from the wonderful VLFeat suite (vlfeat.org) by Andrea Vedaldi and Brian Fulkerson; graph-based image segmentation software by Pedro Felzenszwalb; and the OpenSURF feature detector by Dirk-Jan Kroon, and the Camera Calibration Toolbox by Jean-Yves Bouguet.

As I wrote I became fascinated by the mathematicians, scientists and engineers whose work, hundreds of years ago, underpins the science of robotic vision today. Some of their names have become adjectives like Gaussian, Laplacian or Cartesian or

nouns like Jacobian. They are interesting characters from a distant era when science was a hobby and their day jobs were as doctors, alchemists, gamblers, astrologers, philosophers or mercenaries. In order to know whose shoulders we are standing on I have included small vignettes about the lives of some of these people – a smattering of history as a backstory.

In my own career I have had the good fortune to work with many wonderful people who have inspired and guided me. Long ago at the University of Melbourne John Anderson fired my interest in control and Graham Holmes tried with mixed success to have me “think before I code”. Early on I spent a life-direction-changing ten months working with Richard (Lou) Paul in the GRASP laboratory at the University of Pennsylvania in the period 1988–1989. The genesis of the Toolboxes was my Ph.D. research (1991–1994) and my advisors Malcolm Good (University of Melbourne) and Paul Dunn (CSIRO) asked me good questions and guided my research. Laszlo Nemes (CSIRO) provided great wisdom about life and the ways of organizations, and encouraged me to publish and to open-source my software. Much of my career was spent at CSIRO where I had the privilege and opportunity to work on a diverse range of real robotics and vision projects and to work with a truly talented set of colleagues and friends. Part way through writing the first edition I joined the Queensland University of Technology which made time available to complete that work, and in 2015 sabbatical leave to complete the second.

Many people have helped me in my endeavor and I thank them. I was generously hosted for periods of productive writing at Oxford (both editions) by Paul Newman, and at MIT (first edition) by Daniela Rus. Daniela, Paul and Cédric Pradalier made constructive suggestions and comments on early drafts of that edition. For the second edition I was helped by comments on draft chapters by: Tim Barfoot, Dmitry Bratanov, Duncan Campbell, Donald Dansereau, Tom Drummond, Malcolm Good, Peter Kujala, Obadiah Lam, Jörn Malzahn, Felipe Nascimento Martins, Ajay Pandey, Cédric Pradalier, Dan Richards, Daniela Rus, Sareh Shirazi, Surya Singh, Ryan Smith, Ben Talbot, Dorian Tsai and Ben Upcroft; and assisted with wisdom and content by: François Chaumette, Donald Dansereau, Kevin Lynch, Robert Mahony and Frank Park.

I have tried my hardest to eliminate errors but inevitably some will remain. Please email bug reports to me at rvc@petercorke.com as well as suggestions for improvements and extensions.

Writing the second edition was financially supported by EPSRC Platform Grant EP/M019918/1, QUT Science & Engineering Faculty sabbatical grant, QUT Vice Chancellor’s Excellence Award, QUT Robotics and Autonomous Systems discipline and the ARC Centre of Excellence for Robotic Vision (grant CE140100016).

Over both editions I have enjoyed wonderful support from MathWorks, through their author program, and from Springer. My Springer editor Thomas Ditzinger has been a great supporter of this project and Armin Stasch, with enormous patience and dedication in layout and typesetting, has transformed my untidy ideas into a thing of beauty.

Finally, my deepest thanks are to Phillipa who has supported me and “the book” with grace and patience for a very long time and in many different places – without her this book could never have been written.

Peter Corke
Brisbane,
Queensland
August 2018

Note on the Second Edition

The revision principle was to keep the good (narrative style, code as a first-class citizen, soft plastic cover) and eliminate the bad (errors and missing topics). There were more errors than I would have liked and I thank everybody who submitted errata and suggested improvements.

The first edition was written before I taught in the university classroom or created the MOOCs, which is the inverse of the way books are normally developed. Preparing for teaching gave me insights into better ways to present some topics, particularly around pose representation, robot kinematics and dynamics so the presentation has been adjusted accordingly.

New content includes matrix exponential notation; the basics of screw theory and Lie algebra; inertial navigation; differential steer and omnidirectional mobile robots; a deeper treatment of SLAM systems including scan matching and pose graphs; greater use of MATLAB computer algebra; operational space control; deeper treatment of manipulator dynamics and control; visual SLAM and visual odometry; structured light; bundle adjustment; and light-field cameras.

In the first edition I shied away from Lie algebra, matrix exponentials and twists but I think it's important to cover them. The topic is deeply mathematical and I've tried to steer a middle ground between hardcore algebraic topology and the homogeneous transformation only approach of most other texts, while also staying true to the overall approach of this book.

All MATLAB generated figures have been regenerated to reflect recent improvements to MATLAB graphics and all code examples have been updated as required and tested, and are available as MATLAB Live Scripts.

The second edition of the book is matched by new major releases of my Toolboxes: Robotics Toolbox (release 10) and the Machine Vision Toolbox (release 4). These newer versions of the toolboxes have some minor incompatibilities with previous releases of the toolboxes, and therefore also with the code examples in the first edition of the book.

Note on this Edition

This book is essentially Parts I and IV of the second edition of Robotics, Vision & Control. The previous Part IV has been rearranged and split into two parts, and the content related to robotics and vision-based control has been omitted.

Contents

1	Introduction	1
1.1	About the Book	9
1.1.1	MATLAB Software and the Toolboxes	10
1.1.2	Notation, Conventions and Organization	11
1.1.3	Audience and Prerequisites	11
1.1.4	Learning with the Book	12
1.1.5	Teaching with the Book	12
1.1.6	Outline	13
	Further Reading	14
	Part I Foundations	15
2	Representing Position and Orientation	17
2.1	Working in Two Dimensions (2D)	22
2.1.1	Orientation in 2-Dimensions	23
2.1.2	Pose in 2-Dimensions	26
2.2	Working in Three Dimensions (3D)	31
2.2.1	Orientation in 3-Dimensions	32
2.2.2	Pose in 3-Dimensions	46
2.3	Advanced Topics	49
2.3.1	Normalization	49
2.3.2	Understanding the Exponential Mapping	50
2.3.3	More About Twists	52
2.3.4	Derivative of Pose	55
2.3.5	Transforming Spatial Velocities	56
2.3.6	Incremental Rotation	58
2.3.7	Incremental Rigid-Body Motion	59
2.4	Using the Toolbox	59
2.5	Wrapping Up	61
	Further Reading	63
	Exercises	64
	Part II Computer Vision	67
3	Light and Color	71
3.1	Spectral Representation of Light	71
3.1.1	Absorption	73
3.1.2	Reflectance	74
3.1.3	Luminance	74
3.2	Color	75
3.2.1	The Human Eye	76
3.2.2	Measuring Color	78
3.2.3	Reproducing Colors	79
3.2.4	Chromaticity Space	82
3.2.5	Color Names	84
3.2.6	Other Color and Chromaticity Spaces	85

3.2.7	Transforming between Different Primaries	88
3.2.8	What Is White?	90
3.3	Advanced Topics	90
3.3.1	Color Temperature	90
3.3.2	Color Constancy	91
3.3.3	White Balancing	92
3.3.4	Color Change Due to Absorption	92
3.3.5	Dichromatic Reflectance	94
3.3.6	Gamma	94
3.4	Application: Color Image	96
3.4.1	Comparing Color Spaces	96
3.4.2	Shadow Removal	97
3.5	Wrapping Up	99
	Further Reading	100
	Data Sources	100
	Exercises	101
4	Images and Image Processing	103
4.1	Obtaining an Image	103
4.1.1	Images from Files	103
4.1.2	Images from an Attached Camera	107
4.1.3	Images from a Movie File	109
4.1.4	Images from the Web	110
4.1.5	Images from Maps	111
4.1.6	Images from Code	111
4.2	Image Histograms	113
4.3	Monadic Operations	114
4.4	Dyadic Operations	116
4.5	Spatial Operations	120
4.5.1	Linear Spatial Filtering	120
4.5.2	Template Matching	131
4.5.3	Nonlinear Operations	136
4.6	Mathematical Morphology	137
4.6.1	Noise Removal	140
4.6.2	Boundary Detection	142
4.6.3	Hit or Miss Transform	142
4.6.4	Distance Transform	143
4.7	Shape Changing	145
4.7.1	Cropping	145
4.7.2	Image Resizing	146
4.7.3	Image Pyramids	147
4.7.4	Image Warping	148
4.8	Wrapping Up	151
	Further Reading	151
	Sources of Image Data	153
	MATLAB Notes	153
	General Software Tools	153
	Exercises	154
5	Image Feature Extraction	157
5.1	Region Features	159
5.1.1	Classification	159
5.1.2	Representation	168
5.1.3	Description	171
5.1.4	Summary	181

5.2	Line Features	182
5.2.1	Summary	187
5.3	Point Features	187
5.3.1	Classical Corner Detectors	187
5.3.2	Scale-Space Corner Detectors	193
5.4	Wrapping Up	198
	MATLAB Notes	198
	Further Reading	199
	Exercises	201
	Part III The Geometry of Vision	203
6	Image Formation	207
6.1	Perspective Camera	207
6.1.1	Perspective Projection	207
6.1.2	Modeling a Perspective Camera	210
6.1.3	Discrete Image Plane	212
6.1.4	Camera Matrix	213
6.1.5	Projecting Points	215
6.1.6	Lens Distortion	218
6.2	Camera Calibration	219
6.2.1	Homogeneous Transformation Approach	219
6.2.2	Decomposing the Camera Calibration Matrix	221
6.2.3	Pose Estimation	222
6.2.4	Camera Calibration Toolbox	223
6.3	Wide Field-of-View Imaging	224
6.3.1	Fisheye Lens Camera	225
6.3.2	Catadioptric Camera	228
6.3.3	Spherical Camera	230
6.4	Unified Imaging	232
6.4.1	Mapping Wide-Angle Images to the Sphere	233
6.4.2	Mapping from the Sphere to a Perspective Image	235
6.5	Novel Cameras	236
6.5.1	Multi-Camera Arrays	236
6.5.2	Light-Field Cameras	236
6.6	Advanced Topics	239
6.6.1	Projecting 3D Lines and Quadrics	239
6.6.2	Nonperspective Cameras	241
6.6.3	Camera and Image Motion	242
6.7	Wrapping Up	247
	Further Reading and Resources	248
	Toolbox Notes	249
	Exercises	250
7	Using Multiple Images	253
7.1	Feature Correspondence	254
7.2	Geometry of Multiple Views	258
7.2.1	The Fundamental Matrix	260
7.2.2	The Essential Matrix	262
7.2.3	Estimating the Fundamental Matrix from Real Image Data	264
7.2.4	Planar Homography	268
7.3	Stereo Vision	273
7.3.1	Sparse Stereo	273
7.3.2	Dense Stereo Matching	277
7.3.3	Peak Refinement	283
7.3.4	Cleaning up and Reconstruction	285

7.3.5	3D Texture Mapped Display	288
7.3.6	Anaglyphs	289
7.3.7	Image Rectification	290
7.4	Bundle Adjustment	291
7.5	Point Clouds	297
7.5.1	Fitting a Plane	297
7.5.2	Matching Two Sets of Points	299
7.6	Structured Light	301
7.7	Applications	303
7.7.1	Perspective Correction	303
7.7.2	Mosaicing	306
7.7.3	Image Matching and Retrieval	308
7.7.4	Visual Odometry	314
7.8	Wrapping Up	317
	MATLAB and Toolbox Notes	318
	Further Reading	318
	Resources	322
	Exercises	323
	Appendices	327
A	Installing the Toolboxes	329
B	Linear Algebra Refresher	333
B.1	Vectors	333
B.2	Matrices	334
C	Geometry	341
C.1	Euclidean Geometry	341
C.2	Homogeneous Coordinates	350
C.3	Geometric Transformations	354
D	Lie Groups and Algebras	357
E	Linearization, Jacobians and Hessians	363
F	Solving Systems of Equations	367
F.1	Linear Problems	367
F.2	Nonlinear Problems	368
G	Graphs	377
H	Peak Finding	381
	Bibliography	385
	Index	393
	Index of People	393
	Index of Functions, Classes and Methods	394
	General Index	397

Nomenclature

The notation used in robotic vision varies considerably across books and research papers. The symbols used in this book, and their units where appropriate, are listed below. Some symbols have multiple meanings and their context must be used to disambiguate them.

Notation	Description
x^*	desired value of x
x^+	predicted value of x
$x^\#$	measured, or observed, value of x
\hat{x}	estimated value of x
\bar{x}	mean of x or relative value
$x^{(k)}$	k^{th} element of a time series
\boldsymbol{v}	a vector
$\hat{\boldsymbol{v}}$	a unit-vector parallel to \boldsymbol{v}
$\tilde{\boldsymbol{v}}$	homogeneous representation of vector \boldsymbol{v}
$v[i]$	i^{th} element of vector \boldsymbol{v}
v_x	a component of a vector
\boldsymbol{A}	a matrix
$A[i, j]$	the element (i, j) of \boldsymbol{A}
$A_{i,j}$	the element (i, j) of \boldsymbol{A}
$f(x)$	a function of x
$F_x(x)$	the derivative $\partial f / \partial x$
$F_{xy}(x, y)$	the derivative $\partial^2 f / \partial x \partial y$
$\hat{\boldsymbol{q}}$	unit quaternion, $\hat{\boldsymbol{q}} \in \mathbb{S}^3$
$\mathbf{0}_{m \times n}$	an $m \times n$ matrix of zeros
$\mathbf{1}_{m \times n}$	an $m \times n$ matrix of ones

Symbol	Description	Unit
C	camera matrix, $C \in \mathbb{R}^{3 \times 4}$	
E	illuminance (lux)	lx
f	focal length	m
\mathbf{f}	vector of image features	
\mathbb{H}	the set of all quaternions (H for Hamilton)	
$I_{n \times n}$	$n \times n$ identity matrix	
J	Jacobian matrix	
${}^A J_B$	Jacobian transforming velocities in frame B to frame A	
k, K	constant	
K	camera calibration matrix	
L	luminance (nit)	nt
$N(\mu, \sigma^2)$	a normal (Gaussian) distribution with mean μ and standard deviation σ	
\mathbf{p}	an image plane point, $\mathbf{p} \in \mathbb{R}^2$	
\mathbf{P}	a world point, $\mathbf{P} \in \mathbb{R}^3$	
\mathbb{P}^2	projective space of all 2-D points, a 3-tuple	
\mathbb{P}^3	projective space of all 3-D points, a 4-tuple	
\mathbf{q}	generalized coordinates, configuration $\mathbf{q} \in \mathcal{C}$	m, rad
\mathbf{R}	an orthonormal rotation matrix, $\mathbf{R} \in \text{SO}(2)$ or $\text{SO}(3)$	
\mathbb{R}	set of real numbers	
\mathbb{R}^2	set of all 2-D points	
\mathbb{R}^3	set of all 3-D points	
\mathbb{S}^1	unit circle, set of angles $[0, 2\pi)$	
\mathbb{S}^n	unit sphere embedded in \mathbb{R}^{n+1}	
$\mathfrak{se}(n)$	Lie algebra for $\text{SE}(n)$, an $\mathbb{R}^{(n+1) \times (n+1)}$ augmented skew-symmetric matrix	
$\mathfrak{so}(n)$	Lie algebra for $\text{SO}(n)$, an $\mathbb{R}^{n \times n}$ skew-symmetric matrix	
$\text{SE}(n)$	special Euclidean group, the set of all poses in n dimensions, represented by an $\mathbb{R}^{(n+1) \times (n+1)}$ homogeneous transformation matrix	
$\text{SO}(n)$	special orthogonal group, the set of all orientations in n dimensions, represented by an $\mathbb{R}^{n \times n}$ orthogonal matrix	
S	twist in 3 dimensions, $S \in \mathbb{R}^6$	
t	time	s
T	temperature	K
T	optical transmission	m^{-1}
T	homogeneous transformation, $T \in \text{SE}(2)$ or $\text{SE}(3)$	
${}^A T_B$	homogeneous transform representing frame $\{B\}$ with respect to frame $\{A\}$. If A is not given then assumed relative to world coordinate frame 0. Note that ${}^A T_B = ({}^B T_A)^{-1}$	
u, v	camera image plane coordinates	pixels
u_0, v_0	coordinates of the principal point	pixels
\bar{u}, \bar{v}	normalized image plane coordinates, relative to the principal point	m
v	velocity	m s^{-1}
\mathbf{v}	velocity vector	m s^{-1}
X, Y, Z	Cartesian coordinates	
\bar{x}, \bar{y}	normalized image-plane coordinates	
\mathbb{Z}	set of all integers	
\mathbb{Z}^+	the set of all integers greater than zero	

Symbol	Description	Unit
ϕ	luminous flux (lumens)	lm
Γ	3-angle representation of rotation, $\Gamma \in \mathbb{R}^3$	rad
θ	angle	rad
$\theta_r, \theta_p, \theta_y$	roll pitch yaw angles	rad
λ	wavelength	m
λ	an eigenvalue	
ν	spatial velocity, $\nu = (v_x, v_y, v_z, \omega_x, \omega_y, \omega_z) \in \mathbb{R}^6$	m s^{-1} , rad s^{-1}
ξ	abstract representation of Cartesian pose (pronounced ksi)	
${}^A\xi_B$	abstract representation of relative pose, frame $\{B\}$ with respect to frame $\{A\}$ or rigid-body motion from frame $\{A\}$ to $\{B\}$	
π	mathematic constant	
π	a plane	
ρ_w, ρ_h	pixel width and height	m
σ	standard deviation	
Σ	Lie algebra $\Sigma = [\cdot] \in \mathfrak{se}(3)$	
ω	rotational rate	rad s^{-1}
$\boldsymbol{\omega}$	angular velocity vector	rad s^{-1}
Ω	Lie algebra $\Omega = [\cdot]_{\times} \in \mathfrak{so}(3)$	

Operator	Description	MATLAB
$\ \cdot\ $	norm, or length, of vector: $\mathbb{R}^n \mapsto \mathbb{R}$	<code>norm</code> , <code>.norm</code>
$\mathbf{v}_1 \cdot \mathbf{v}_2$	dot, or inner, product, also $\mathbf{v}_1^T \mathbf{v}_2$: $\mathbb{R}^n \times \mathbb{R}^n \mapsto \mathbb{R}$	<code>dot</code>
$\mathbf{v}_1 \times \mathbf{v}_2$	cross, or vector, product: $\mathbb{R}^n \times \mathbb{R}^n \mapsto \mathbb{R}^n$	<code>cross</code>
\mathbf{A}^{-1}	inverse of \mathbf{A} : $\mathbb{R}^{n \times n} \mapsto \mathbb{R}^{n \times n}$	<code>inv</code>
\mathbf{A}^+	pseudo-inverse of \mathbf{A} : $\mathbb{R}^{n \times m} \mapsto \mathbb{R}^{m \times n}$	<code>pinv</code>
\mathbf{A}^*	adjugate of $\mathbf{A} \mapsto \det(\mathbf{A})\mathbf{A}^{-1}$, $\mathbb{R}^{n \times n} \mapsto \mathbb{R}^{n \times n}$	
\mathbf{A}^T	transpose of \mathbf{A} : $\mathbb{R}^{n \times m} \mapsto \mathbb{R}^{m \times n}$	<code>'</code>
\mathbf{A}^{-T}	transpose of inverse $\mathbf{A} \mapsto (\mathbf{A}^T)^{-1} = (\mathbf{A}^{-1})^T$, $\mathbb{R}^{n \times n} \mapsto \mathbb{R}^{n \times n}$	
\bullet	transform a point (coordinate vector) by a relative pose: $\mathbf{SE}(n) \times \mathbb{R}^n \mapsto \mathbb{R}^n$	<code>*</code>
\oplus	composition: $\mathbf{S}_E^0(n) \times \mathbf{S}_E^0(n) \mapsto \mathbf{S}_E^0(n)$	<code>*</code>
\ominus	composition with inverse: $\mathbf{S}_E^0(n) \times \mathbf{S}_E^0(n) \mapsto \mathbf{S}_E^0(n)$	<code>/</code>
\ominus	unary inverse: $\mathbf{S}_E^0(n) \mapsto \mathbf{S}_E^0(n)$	<code>.inv</code>
$\Delta(\cdot)$	maps incremental pose change to differential motion: $\mathbf{SE}(3) \mapsto \mathbb{R}^6$	<code>tr2delta</code>
$\Delta^{-1}(\cdot)$	maps differential motion to incremental pose change: $\mathbb{R}^6 \mapsto \mathbf{SE}(3)$	<code>delta2tr</code>
$\mathcal{R}_i(\theta)$	pure rotation about axis i : $\mathbb{R} \mapsto \mathbf{SE}(3)$	<code>SE3.rotx y z</code>
$\mathcal{R}(\omega)$	pure rotation by $\ \omega\ $ about ω : $\mathbb{R}^3 \mapsto \mathbf{SE}(3)$	<code>SE3.angvec</code>
$\mathcal{T}_i(d)$	pure translation along axis i : $\mathbb{R} \mapsto \mathbf{SE}(2), \mathbf{SE}(3)$	<code>SE2, SE3</code>
$\mathcal{T}(\mathbf{t})$	pure translation by vector: $\mathbb{R}^n \mapsto \mathbf{SE}(n)$	<code>SE2, SE3</code>
$[\cdot]_t$	translational component of pose: $\mathbf{SE}(n) \mapsto \mathbb{R}^n$	<code>.t</code>
$[\cdot]_R$	rotational component of pose: $\mathbf{SE}(n) \mapsto \mathbb{R}^{n \times n}$	<code>.R</code>
$[\cdot]_\times$	skew-symmetric matrix: $\mathbb{R} \mapsto \mathbf{so}(2), \mathbb{R}^3 \mapsto \mathbf{so}(3)$	<code>skew</code>
$\vee_\times(\cdot)$	<i>unpack</i> skew-symmetric matrix: $\mathbf{so}(2) \mapsto \mathbb{R}, \mathbf{so}(3) \mapsto \mathbb{R}^3$	<code>vex</code>
$[\cdot]$	augmented skew-symmetric matrix: $\mathbb{R}^3 \mapsto \mathbf{se}(2), \mathbb{R}^6 \mapsto \mathbf{se}(3)$	<code>skewa</code>
$\vee(\cdot)$	<i>unpack</i> augmented skew-symmetric matrix: $\mathbf{se}(2) \mapsto \mathbb{R}^3, \mathbf{se}(3) \mapsto \mathbb{R}^6$	<code>vexa</code>
$\text{Ad}(\cdot)$	adjoint representation: $\mathbf{SE}(3) \mapsto \mathbb{R}^{6 \times 6}$	<code>.Ad</code>
$\text{ad}(\cdot)$	logarithm of adjoint representation: $\mathbf{SE}(3) \mapsto \mathbb{R}^{6 \times 6}$	<code>.ad</code>
\circ	quaternion (Hamiltonian) multiplication: $\mathbb{H} \times \mathbb{H} \mapsto \mathbb{H}$	<code>*</code>
$\hat{\mathbf{v}}$	pure quaternion: $\mathbb{R}^3 \mapsto \mathbb{H}$	<code>Quaternion.pure</code>
\sim	equivalence of representations	
\simeq	homogeneous coordinate equivalence	
\ominus	smallest angular difference between two angles on a circle: $\mathbb{S}^1 \times \mathbb{S}^1 \mapsto \mathbb{R}$	<code>angdiff</code>
$\mathcal{P}(\cdot)$	camera projection function: $\mathbb{R}^3 \mapsto \mathbb{R}^2$	<code>.project</code>
$*$	convolution	<code>iconv</code>
\otimes	correlation	
\equiv	colormetric equivalence	
\oplus	morphological dilation	
\ominus	morphological erosion	
\circ	morphological opening	
\bullet	morphological closing	
$\{F\}$	coordinate frame F	
$[a, b]$	interval a to b inclusive	
(a, b)	interval a to b exclusive, not including a or b	
$[a, b)$	interval a to b , not including b	
$(a, b]$	interval a to b , not including a	

MATLAB® Toolbox Conventions

- A Cartesian coordinate, a point, is expressed as a column vector.
- A set of points is expressed as a matrix with columns representing the coordinates of individual points.
- A rectangular region is defined by its two opposite corners as a 2×2 matrix $[x_{\min} \ x_{\max}; y_{\min} \ y_{\max}]$.
- Time series data is expressed as a matrix with rows representing time steps.
- A MATLAB matrix has subscripts (i, j) which represent row and column respectively. Image coordinates are written (u, v) so an image represented by a matrix I is indexed as $I(v, u)$.
- Matrices with three or more dimensions are frequently used:
 - A color image has 3 dimensions: row, column, color plane.
 - A greyscale image sequence has 3 dimensions: row, column, index.
 - A color image sequence has 4 dimensions: row, column, color plane, index.

Common Abbreviations

2D	2-dimensional
3D	3-dimensional
n -tuple	A group of n numbers, it can represent a point of a vector